

Iterative Multiplicative Filters for Data Labeling

Ronny Bergmann*, Jan Henrik Fitschen*, Johannes Persch*,
and Gabriele Steidl*

September 14, 2016

Abstract

Based on an idea in [4] we propose a new iterative multiplicative filtering algorithm for label assignment matrices which can be used for the supervised partitioning of data. Starting with a row-normalized matrix containing the averaged distances between prior features and the observed ones the method assigns in a very efficient way labels to the data. We interpret the algorithm as a gradient ascent method with respect to a certain function on the product manifold of positive numbers followed by a reprojection onto a subset of the probability simplex consisting of vectors whose components are bounded away from zero by a small constant. While such boundedness away from zero is necessary to avoid an arithmetic underflow, our convergence results imply that they are also necessary for theoretical reasons. Numerical examples show that the proposed simple and fast algorithm leads to very good results. In particular we apply the method for the partitioning of manifold-valued images.

1 Introduction

Data labeling is a basic problem which appears in various applications. In particular it can be used for image partitioning and segmentation, which is an important preprocessing step for many state-of-the-art algorithms in high-level computer vision. It addresses the task to assign labels from a finite set to the image points in a meaningful way. Thereby the number of labels K is much smaller than the image dimension n . Fig. 1 illustrates a typical labeling result which leads to the segmentation of a texture image.

As a *first ingredient* for a labeling procedure features of the pixels together with a similarity measure for these features is needed. In this paper we only deal with supervised labeling where prototypical features for each label are available. While for natural color images in the RGB color space often the ℓ_p -distance, $p \in \{1, 2\}$, on \mathbb{R}^3 is sufficient, this choice is clearly not appropriate when dealing with manifold-valued images as, e.g., in DT-MRI, where the image pixels are positive definite 3×3 matrices or in Electron Backscattered Diffraction, where the pixels lie in a quotient manifold of the special orthogonal group. Also for the texture image in Fig. 1 we have to deal with more appropriate features, actually means and covariance matrices associated to the image pixels as described in the numerical part in Section 4,

*Department of Mathematics, Technische Universität Kaiserslautern, Postfach 3049, 67653 Kaiserslautern, Germany, {bergmann, fitschen, persch, steidl}@mathematik.uni-kl.de.




Figure 1. Labeling of textures based on means and covariance matrices obtained from vectors containing the gray value and the first and second order differences in horizontal and vertical direction. Different labels are visualized by different colors in the right image.

paragraph on „symmetric positive definite matrices”, second example. In other words, often the features can be considered to lie on a manifold and the distance on the manifold provides a good similarity measure. We remark that there is a large amount of literature on the collection of appropriate image features and a detailed description is beyond the scope of this paper. However, taking just features of single pixels and their distances to prototypical label features into account would lead to a high sensibility of the method to errors. To get robust labeling procedures and to respect moreover the geometry of the image the local or nonlocal neighborhood of the pixels should be involved as a *second ingredient* into the labeling procedure. A first idea would be a two step method which applies a neighborhood filtering procedure directly on the image features followed by a label assignment via comparison of distances to the prototypical label features. The neighborhood filters may be local or nonlocal [11, 19, 23] and linear or nonlinear [48]. Indeed, such nonlinear methods may also comprise nonlinear diffusion filter [10, 51] or variational restoration models inspired by the Mumford-Shah functional [13, 38]. However, such methods become rather complicated if the features lie for example in a non flat manifold. Then already the simple local, linear averaging method calls for the computation of Karcher means which requires itself an iterative procedure. Note that recently nonlocal patch-based methods were extensively applied for the denoising of manifold-valued images in [34]. In this paper, we will apply the neighborhood filtering not directly on the image features, but on the so-called label assignment vectors. The label assignment vectors $W_i := (W_{i,k})_{k=1}^K$, $i = 1, \dots, n$, contain the probabilities $W_{i,k}$ that image pixel i belongs to label k . Clearly, since we are dealing with probabilities, we have that W_i belongs to the probability simplex

$$\Delta_K := \left\{ x \in \mathbb{R}^K : \sum_{k=1}^K x_k = 1, x \geq 0 \right\}. \quad (1)$$

Hence, our neighborhood filtering process is independent of the considered feature space. The idea to handle neighborhood relations via the label assignment vectors can also be found in so-called relaxation labeling approaches [29, 41, 45]. However these methods do not use a multiplicative filtering process.

Besides the above mentioned methods there exists a huge number of variational labeling approaches and no single technique works best for all cases. The models can roughly be divided into continuous and discrete ones. Recent continuous multilabel models as those in [8, 14, 36] consider a convex functional which consists of a data term that contains the distances between the pixel features and prototypical features weighted by the label assignment vectors and a regularization term which enforces certain neighborhood relations on the label assignment vectors. A frequent choice are isotropic total variation regularizers [46] and their relatives as wavelet-like terms [25]. The global minimizers can be found by primal dual first order methods [12, 15]. In [16] it was shown that the relaxed two label method is tight, a promising result which is unfortunately no longer true for multilabel tasks [35]. Discrete models, in particular relaxed linear programs, are often the method of choice in computer vision since they appear to be more efficient for huge sets of high dimensional data. Moreover they are in general tighter than continuous models, see exemplary [31, 32]. For a comprehensive study of discrete energy minimizing methods also in the context of image segmentation we refer to [30] which significantly expands the scope of methods involved in the earlier comparative study [47]. The above mentioned continuous and discrete models are non-smooth and convex. We want to mention that there exist also non convex approaches as, e.g., those in [26, 27, 40, 50, 53] which always have to cope with local minima.

Recently, Åström, Petra, Schmitzer, and Schnörr [4] suggested an interesting supervised geometric approach to the labeling problem, see also [3, 5] for the corresponding conference papers. The objective function to maximize is defined on the open manifold of stochastic matrices with the Fisher-Rao metric and a maximizing algorithm via the corresponding Riemannian gradient ascent flow is considered. In the numerical part the authors apply several simplifications, in particular lifting maps to cope with the non completeness of the manifold. Finally this leads to an iterative numerical procedure. Unlike the continuous Riemannian gradient flow that is shown in [4] to converge to edges of the simplex, the authors merely showed that the simplified numerical scheme closely approximates this flow, but did not prove its convergence. In this paper we discuss the convergence of the resulting numerical scheme.

Since the approach in [4] may be hard to read for people not trained in differential geometry we show in Section 3 how their final iteration scheme can be rewritten such that we arrive at an easily understandable algorithm. A slight modification of this algorithm leads to our (even simpler) algorithm reported in the next section.

We start with this algorithm since we want to introduce it from a user-friendly point of view as a multiplicative filtering process of label assignment vectors. We initialize our iterations with a label assignment matrix containing label assignment vectors of averaged distances between the prior features and the observed ones. Matrices of this kind are common kernel matrices in nonlinear filter [11, 48], patch- and graph Laplacian based methods, see, e.g., [9, 20]. In particular the data and their priors may lie in any metric space, which makes the method highly flexible. Then these label assignment vectors are iterated in a multiplicative way, where we have to take care that the components stay within the probability simplex and a (small) fixed distance $\varepsilon > 0$ away from zero. We analyze the convergence of the algorithm if such an ε is not fixed. In exact arithmetic it appears that in most relevant cases the algorithm would converge to a trivial label assignment matrix. This is not the case if we demand that the components are bounded away from zero by a small distance. Such condition is used in the numerical part anyway to avoid an arithmetic underflow. However,

the analysis shows that there is also a theoretical reason for bounding the components away from zero. Indeed, in all numerical examples we observed convergence of the algorithm to the vertices of the ε probability simplex $\Delta_{K,\varepsilon} := \{x \in \mathbb{R}^K : \sum_{k=1}^K x_k = 1, x_i \geq \varepsilon, i = 1, \dots, K\}$. We reinterpret our algorithm as gradient ascent step of a certain function on the product manifold of positive numbers followed by a reprojection onto the ε probability simplex. The projection is done with respect to the Kullback-Leibler distance. We provide various numerical examples.

The outline of this paper is as follows: In Section 2 we introduce and discuss our iterative multiplicative filtering algorithm: the simple algorithm is presented in Subsection 2.1, followed by its convergence analysis in the case that the components of the label assignment vectors are not bounded away from zero in Subsection 2.2. We interpret the algorithm as a gradient ascent step of a certain function followed by a Kullback-Leibler projection onto the ε probability simplex in Subsection 2.3. The relation to paper [4] together is detailed in Section 3. Numerical results demonstrate the very good performance of our algorithm in Section 4. In particular we apply the method for the partitioning of manifold-valued images. The paper finishes with conclusions in Section 5.

2 Multiplicative Filtering of Label Assignment Matrices

2.1 Basic Algorithm

In this subsection we introduce the supervised labeling algorithm. As outlined in the introduction two basic ingredients are needed, namely i) feature vectors (prototypical ones for the labels of each pixel) together with a distance measure between them, and ii) neighborhood relations between pixels usually given by certain weights.

Let I_n denote the $n \times n$ identity matrix, $\mathbf{1}_{n,m}$ the $n \times m$ matrix with all components 1, and $\mathbf{e}_k \in \mathbb{R}^K$ the k -th unit vector in \mathbb{R}^K . By $\|x\|_1 := |x_1| + \dots + |x_K|$ we denote the ℓ_1 norm and by $\langle x, y \rangle := x_1 y_1 + \dots + x_K y_K$ the inner product of vectors $x, y \in \mathbb{R}^K$. Likewise, if $X, Y \in \mathbb{R}^{n,K}$ are matrices, we use $\langle X, Y \rangle = \sum_{j=1}^n \sum_{k=1}^K X_{j,k} Y_{j,k}$. Let $(\mathcal{M}, \text{dist})$ be a metric space. In our applications \mathcal{M} will be \mathbb{R}^d or a non-flat Riemannian manifold.

We are interested in partitioning n features

$$f_i \in \mathcal{M}, \quad i = 1, \dots, n,$$

into $K \ll n$ labels, where we suppose that we are given K prior features

$$f_k^* \in \mathcal{M}, \quad k = 1, \dots, K.$$

In other words, we consider supervised labeling. Then

$$\mathcal{D} := (d_{i,k})_{i=1,k=1}^{n,K} = (\mathcal{D}_1 \dots \mathcal{D}_n)^T, \quad d_{i,k} := \text{dist}(f_i, f_k^*), \quad (2)$$

is the distance matrix of the features to the priors. Throughout this paper exponentials and logarithms of matrices are meant componentwise. Further, $x \circ y$ denotes the componentwise product of matrices. To each pixel $i \in \{1, \dots, n\}$ we assign a neighborhood $\mathcal{N}_\alpha(i) \subseteq \{1, \dots, n\}$ and set

$$A_i := \exp \left(- \sum_{j \in \mathcal{N}_\alpha(i)} \alpha_{i,j} \mathcal{D}_j \right) / \left\| \exp \left(- \sum_{j \in \mathcal{N}_\alpha(i)} \alpha_{i,j} \mathcal{D}_j \right) \right\|_1 \in \mathbb{R}_{>0}^K \quad (3)$$

and

$$A := (A_1 \dots A_n)^T \in \mathbb{R}_{\geq 0}^{n,K},$$

where $\alpha_{i,j} \in (0, 1)$ with $\sum_{j \in \mathcal{N}_\alpha(i)} \alpha_{i,j} = 1$. We can extend the definition of the weights $\alpha_{i,j}$ to all $j = 1, \dots, n$ by setting $\alpha_{i,j} = 0$ if $j \notin \mathcal{N}_\alpha(i)$.

We will work on *label assignment matrices* with non-negative entries

$$W := (W_1 \dots W_n)^T \in \mathbb{R}_{\geq 0}^{n,K},$$

whose rows sum up to 1. In other words, W consists of *label assignment vectors* $W_i \in \Delta_K$, where Δ_K denotes the probability simplex (1). In particular, $W_i^{(0)} := A_i \in \Delta_K$ will serve as initialization for our algorithm. Then we apply a multiplicative neighborhood averaging to the columns of W , i.e., separately for each label $k \in \{1, \dots, K\}$. To have full flexibility, we may consider other neighborhoods \mathcal{N}_ρ and weights $\rho_{i,j} > 0$ with $\sum_{j \in \mathcal{N}_\rho(i)} \rho_{i,j} = 1$, $i = 1, \dots, n$ for this process. We assume that the weights in Algorithm 1 fulfill

$$\rho_{i,j} = \rho_{j,i},$$

i.e., mutual neighbors are weighted in the same way and extend their definition to all indices in $\{1, \dots, n\}$ by setting $\rho_{i,j} := 0$ if $j \notin \mathcal{N}_\rho(i)$. In summary we propose the Algorithm 1:

Algorithm 1 Multiplicative Filtering

Input: $W^{(0)} := A$, $\rho_{i,j} \in [0, 1]$ with $\sum_{j \in \mathcal{N}(i)} \rho_{i,j} = 1$, $i = 1, \dots, n$, $\varepsilon := 10^{-10}$

$r = 0$;

repeat

for $i = 1, \dots, n$ **do**

$$1. U_i^{(r+1)} = W_i^{(r)} \circ \prod_{j \in \mathcal{N}_\rho(i)} (W_j^{(r)})^{\rho_{i,j}}$$

$$2. V_i^{(r+1)} = U_i^{(r+1)} / \|U_i^{(r+1)}\|_1$$

3. Set $W_i^{(r+1)} := V_i^{(r+1)}$ and proceed

$$\mathcal{I} = \{k : W_{i,k}^{(r+1)} \leq \varepsilon\}$$

while $\exists k \in \{1, \dots, K\} : W_{i,k}^{(r+1)} < \varepsilon$ **do**

$$\mathcal{I} = \mathcal{I} \cup \arg \min_{k=1, \dots, K} W_{i,k}^{(r+1)}$$

$$\tau_{\mathcal{I}} = \frac{1 - |\mathcal{I}| \varepsilon}{1 - \sum_{k \in \mathcal{I}} W_{i,k}^{(r+1)}}$$

$$W_{i,k}^{(r+1)} = \begin{cases} \varepsilon & k \in \mathcal{I} \\ \tau_{\mathcal{I}} W_{i,k}^{(r+1)} & \text{otherwise} \end{cases}$$

$r \rightarrow r + 1$;

until a stopping criterion is reached

When the algorithm stops with $W_i^{(R)}$, we assign the label

$$\arg \max_{k \in \{1, \dots, K\}} W_{i,k}^{(R)} \tag{4}$$

to the i -th pixel. Let us comment the steps of the algorithm:

Step 1. Here $\prod_{j \in \mathcal{N}_\rho(i)} (W_j)^{\rho_{i,j}}$ is the weighted geometric mean of the label assignment vectors W_j in the neighborhood of i . In particular we have for $\rho_{i,j} := 1/|\mathcal{N}_\rho(i)|$ that

$$\prod_{j \in \mathcal{N}_\rho(i)} (W_{j,k})^{\rho_{i,j}} = \left(\prod_{j \in \mathcal{N}_\rho(i)} W_{j,k} \right)^{\frac{1}{|\mathcal{N}_\rho(i)|}}, \quad k = 1, \dots, K.$$

Taking the logarithm we obtain

$$\log U^{(r)} = \log W^{(r)} + P \log W^{(r)}$$

with the weight matrix

$$P := (\rho_{i,j})_{i,j=1}^n. \quad (5)$$

Hence Step 1 can be rewritten as

$$U^{(r+1)} = W^{(r)} \circ \exp\left(P \log W^{(r)}\right).$$

Step 2. Here we ensure that the rows of $V^{(r+1)}$ lie in the probability simplex Δ_K .

Step 3. This step basically avoids that the entries of $W^{(r)}$ become too small (actually smaller than $\varepsilon = 10^{-10}$). Very small values do also not appear if we set

$$W_i^{(r+1)} := V_i^{(r+1)} + \mathbf{1}_K \left(\varepsilon - \min_k \{V_{i,k}^{(r+1)}\} \right)$$

followed by a normalization. Indeed, in numerical tests this update instead of Step 3 works similar. We have chosen Step 3 since together with Step 2 it has a nice interpretation as Kullback-Leibler projection onto the part of the probability simplex with entries larger or equal to ε . However, we will see in the next subsection that Step 3 is not just numerical cosmetics, but in general absolutely necessary to avoid that the iterates $W_i^{(r)}$ converge to the same vertex of the simplex for all $i = 1, \dots, n$ which would result in a trivial labeling.

By Step 1, the weight matrix P plays a crucial role. By definition this matrix is symmetric and stochastic, i.e., it has nonnegative entries and $P\mathbf{1}_n = \mathbf{1}_n$. We finish this subsection by collecting some well-known properties of symmetric stochastic matrices, see, e.g., [22, 28, 43, 52].

Lemma 2.1. *Let $P \in \mathbb{R}_{\geq 0}^{n,n}$ be a symmetric stochastic matrix. Then the following holds true:*

i) *The matrix P has an eigenvalue decomposition of the form*

$$P = Q\Lambda Q^T, \quad \Lambda := \text{diag}(\lambda_1, \dots, \lambda_n), \quad (6)$$

where $\lambda_1 = 1$, $|\lambda_i| \leq 1$, $i = 2, \dots, n$ and

$$Q := (q_1 \dots q_n)^T$$

is an orthogonal matrix with first column $(q_{j,1})_{j=1}^n = \mathbf{1}_n / \sqrt{n}$.

- ii) If P has positive diagonal entries, then -1 is not an eigenvalue.
- iii) If P is irreducible, i.e., the associated matrix graph is connected, then 1 is a single eigenvalue.
- iv) If P has positive diagonal entries and $\lambda_1 = 1$ is a single eigenvalue, then P^r converges as $r \rightarrow \infty$ to the constant matrix $\frac{1}{n} \mathbf{1}_{n,n}$.

2.2 Convergence Analysis for $\varepsilon = 0$

In this subsection, we analyze the convergence of Algorithm 1 without Step 3, i.e., we use the same initialization as in the algorithm, but compute just

$$\begin{aligned} 1. \quad U_i^{(r+1)} &= W_i^{(r)} \circ \prod_{j \in \mathcal{N}_\rho(i)} (W_j^{(r)})^{\rho_{i,j}}, \\ 2. \quad W_i^{(r+1)} &= U_i^{(r+1)} / \|U_i^{(r+1)}\|_1. \end{aligned} \tag{7}$$

In other words we consider the case $\varepsilon = 0$. We will see that the third step is indeed an essential.

We start with the following observation.

Lemma 2.2. *The r -th iterate in (7) is given by*

$$W^{(r)} = \left(\frac{\exp w_1^{(r)}}{\|\exp w_1^{(r)}\|_1}, \dots, \frac{\exp w_n^{(r)}}{\|\exp w_n^{(r)}\|_1} \right)^T,$$

with

$$w^{(r)} = (w_1^{(r)} \dots w_n^{(r)})^T := (I_n + P)^r a, \quad a := \log A$$

and the weight matrix P in (5).

Proof. First note that we obtain the same iterates $W_i^{(r)}$ if we skip the normalization Step 2 in the intermediate iterations and normalize only in step r . Therefore we consider the sequence with the same starting point $\tilde{W}^{(0)} := W^{(0)}$ and iterations

$$\tilde{W}_i^{(r+1)} := \tilde{W}_i^{(r)} \circ \prod_{j \in \mathcal{N}(i)} (\tilde{W}_j^{(r)})^{\rho_{i,j}}.$$

Taking the logarithm and setting $w_i^{(r)} := \log \tilde{W}_i^{(r)}$ the iteration becomes

$$w_i^{(r+1)} = w_i^{(r)} + \sum_{j \in \mathcal{N}(i)} \rho_{i,j} w_j^{(r)}, \tag{8}$$

where $w^{(0)} = \log W^{(0)} = a$. With $w^{(r)} := (w_1^{(r)} \dots w_n^{(r)})^T$ this can be rewritten as

$$\begin{aligned} w^{(r+1)} &= (I_n + P)w^{(r)} \\ &= (I_n + P)^{r+1}w^{(0)}. \end{aligned} \tag{9}$$

Then $\tilde{W}^{(r)} = \exp w^{(r)}$ and row normalization yields the assertion. \square

The following remark gives a first intuition about the convergence behavior of $W^{(r)}$.

Remark 2.3. By (9) we have

$$w^{(r)} = 2^r \left(\frac{1}{2}(I_n + P) \right)^r w^{(0)} = 2^r v^{(r)}.$$

From the definition of P we know that $\frac{1}{2}(I_n + P)$ is a symmetric stochastic matrix with positive diagonal entries. Assuming that 1 is a single eigenvalue of this matrix, we conclude by Lemma 2.1 that $v^{(r)}$ converges to the matrix $\frac{1}{n}\mathbf{1}_n\mathbf{1}_n^T w^{(0)}$ with constant columns as $r \rightarrow \infty$. However, this does in general not imply that the columns of $W^{(r)}$ also converge to a constant vector for $r \rightarrow \infty$ as the following example shows: We have

$$v^{(r)} := \begin{pmatrix} \frac{1}{2} + \frac{1}{2^r} & \frac{1}{2} + (\frac{2}{3})^r \\ \frac{1}{2} + (\frac{2}{3})^r & \frac{1}{2} + \frac{1}{2^r} \end{pmatrix} \rightarrow \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} =: v \quad \text{as } r \rightarrow \infty.$$

Multiplying with 2^r , taking the exponential and normalizing the rows leads in the first column of $v^{(r)}$ to

$$\begin{aligned} \frac{\exp\left(2^r \left(\frac{1}{2} + \frac{1}{2^r}\right)\right)}{\exp\left(2^r\left(\frac{1}{2} + \frac{1}{2^r}\right)\right) + \exp\left(2^r\left(\frac{1}{2} + (\frac{2}{3})^r\right)\right)} &= \frac{\exp\left(\left(2^{r-1} + 1\right)\right)}{\exp\left(2^{r-1} + 1\right) + \exp\left(2^{r-1} + (\frac{4}{3})^r\right)} \\ &= \frac{\exp\left(\left(\frac{1}{(\frac{4}{3})^r}\right)\right)}{\exp(1) + \exp\left((\frac{4}{3})^r\right)} \end{aligned}$$

so that

$$\lim_{r \rightarrow \infty} \frac{\exp(2^r v^{(r)})}{\exp(2^r v^{(r)})_{1,1} + \exp(2^r v^{(r)})_{1,2}} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

The general convergence result is given next.

Theorem 2.4. i) The sequence $\{W_i^{(r)}\}_r$ generated by (7) converges.

ii) Let P have the eigenvalue decomposition (6), where Λ contains the eigenvalues in descending order. Assume that there are M pairwise distinct eigenvalues λ_{n_s} with multiplicity κ_{n_s} , $s = 1, \dots, M$. Let \hat{s} be the largest index such that $\lambda_{n_{\hat{s}}} > 0$. Set

$$\hat{a} := Q^T \log A, \quad \hat{a} = (\hat{a}_1 \dots \hat{a}_K) \in \mathbb{R}^{n,K}.$$

Then, for every $i \in \{1, \dots, n\}$, the sequence $\{W_i^{(r)}\}_r$ converges to a unit vector if and only if A fulfills the following property

(PI): If for an $i \in \{1, \dots, n\}$, there exist $k, l \in \{1, \dots, K\}$, $l \neq k$ such that

$$c_{l,k}(s) := \sum_{j=n_s}^{n_s+\kappa_{n_s}-1} q_{i,j}(\hat{a}_{j,l} - \hat{a}_{j,k}) = 0 \quad \text{for all } s = 1, \dots, \hat{s},$$

then there exists $\tilde{l} \in \{1, \dots, K\}$ and $\tilde{s} = \tilde{s}(k, \tilde{l}) \in \{1, \dots, \hat{s}\}$ so that

$$c_{\tilde{l},k}(\tilde{s}) > 0 \quad \text{and} \quad c_{\tilde{l},k}(s) = 0 \quad \text{for all } s = 1, \dots, \tilde{s} - 1.$$

Proof. i) By (9) and (6) we obtain

$$w^{(r)} = (I_n + P)^r w^{(0)} = Q(I_n + \Lambda)^r \hat{a} = Q D^r \hat{a}, \quad (10)$$

where $D := \text{diag}(\mu_1, \dots, \mu_n)$ and $\mu_j = 1 + \lambda_j$, $j = 1, \dots, n$. By Lemma 2.1 and Gershgorin's circle theorem we know that $\mu_j \in (0, 2]$, $j = 1, \dots, n$ and $\mu_j \in (1, 2]$ if and only if $j \leq n_{\hat{s}}$. Taking the exponential and normalizing the rows we get for the i -th row

$$\begin{aligned} W_i^{(r)} &= \frac{1}{\sum_{k=1}^K \exp(q_i^T D^r \hat{a}_k)} (\exp(q_i^T D^r \hat{a}_1) \dots \exp(q_i^T D^r \hat{a}_K))^T \\ &= \left(\frac{1}{1 + \frac{\exp(q_i^T D^r \hat{a}_2)}{\exp(q_i^T D^r \hat{a}_1)} + \dots + \frac{\exp(q_i^T D^r \hat{a}_K)}{\exp(q_i^T D^r \hat{a}_1)}} \dots \frac{1}{\frac{\exp(q_i^T D^r \hat{a}_1)}{\exp(q_i^T D^r \hat{a}_K)} + \dots + \frac{\exp(q_i^T D^r \hat{a}_{K-1})}{\exp(q_i^T D^r \hat{a}_K)} + 1} \right)^T \\ &= \left(\frac{1}{1 + G_{2,1}(r) + \dots + G_{K,1}(r)} \dots \frac{1}{G_{1,K}(r) + \dots + G_{K-1,K}(r) + 1} \right)^T, \end{aligned} \quad (11)$$

where

$$G_{l,k}(r) := \frac{\exp(q_i^T D^r \hat{a}_l)}{\exp(q_i^T D^r \hat{a}_k)} = \prod_{j=1}^n \exp(q_{i,j} \mu_j^r (\hat{a}_{j,l} - \hat{a}_{j,k})).$$

Taking the logarithm results in

$$\begin{aligned} g_{l,k}(r) &:= \log(G_{l,k}(r)) = \sum_{j=1}^n q_{i,j}(\hat{a}_{j,l} - \hat{a}_{j,k}) \mu_j^r = \sum_{s=1}^M c_{l,k}(s) \mu_{n_s}^r \\ &= \sum_{s=1}^{\hat{s}} c_{l,k}(s) \mu_{n_s}^r + \sum_{s=\hat{s}+1}^M c_{l,k}(s) \mu_{n_s}^r. \end{aligned}$$

The first sum contains the eigenvalues larger than 1 and the second one those smaller or equal to 1. If it exists, let $\tilde{s} \in \{1, \dots, \hat{s}\}$ be the smallest index such that $c_{l,k}(\tilde{s}) \neq 0$. Then it follows

$$\lim_{r \rightarrow \infty} g_{l,k}(r) = \begin{cases} \text{const} & \text{if } c_{l,k}(s) = 0 \text{ for all } s = 1, \dots, \hat{s}, \\ \text{sgn}(c_{l,k}(\tilde{s})) \infty & \text{if } c_{l,k}(s) = 0 \text{ for all } s = 1, \dots, \tilde{s} - 1, \end{cases}$$

and consequently

$$\lim_{r \rightarrow \infty} G_{l,k}(r) = \begin{cases} \exp(\text{const}) & \text{if } c_{l,k}(s) = 0 \text{ for all } s = 1, \dots, \hat{s}, \\ +\infty & \text{if } c_{l,k}(\tilde{s}) > 0, \\ 0 & \text{if } c_{l,k}(\tilde{s}) < 0. \end{cases} \quad (12)$$

Hence, by (11), we see that $\{W_i^{(r)}\}_r$ converges as $r \rightarrow \infty$.

ii) It remains to examine under which conditions it converges to a unit vector. Consider a fixed $k \in \{1, \dots, K\}$. We distinguish two cases.

1. If for all $l \in \{1, \dots, K\} \setminus \{k\}$ there exists an $s = s(l) \in \{1, \dots, \hat{s}\}$ such that $c_{l,k}(s) \neq 0$, then we see by (12) that $\lim_{r \rightarrow \infty} W_{i,k}^{(r)} \in \{0, 1\}$. Since, by construction the components of $W_i^{(r)}$ sum up to 1, we conclude that $W_i^{(r)}$ converges to a unit vector.

2. If there exists $l \in \{1, \dots, K\} \setminus \{k\}$ such that $c_{k,l}(s) = 0$ for all $s = 1, \dots, \hat{s}$, then, by (12), the component $W_{i,k}^{(r)} = (1 + G_{l,k}(r) + \sum_{m \neq l,k} G_{m,k}(r))^{-1}$ is strictly smaller than 1 and becomes 0 if and only if there exists $\tilde{l} \in \{1, \dots, K\}$ and $\tilde{s} = \tilde{s}(k, \tilde{l}) \in \{1, \dots, \hat{s}\}$ so that $c_{\tilde{l},k}(\tilde{s}) > 0$ and $c_{\tilde{l},k}(s) = 0$ for all $s = 1, \dots, \tilde{s} - 1$. This is exactly condition (PI) and we are done. \square

Let us have a closer look at the property (PI).

If $\mathcal{N}(i) = \{i\}$ for all $i \in \{1, \dots, n\}$ which means that we do not take the neighborhoods of the points into account, then $P = I$ and property (PI) is fulfilled if for every $i \in \{1, \dots, n\}$ the vector $(\text{dist}(f_i, f_k^*))_{k=1}^K$ has a unique smallest element. Then the algorithm gives the same assignment as just $\arg \min_k \text{dist}(f_i, f_k^*)$ which is reasonable.

Things change completely if we take care of the point neighborhoods. First we give an equivalent condition to (PI).

Remark 2.5. *Having a look at part ii) of the above proof we realize by (12) and (11) that $\{W_i^{(r)}\}_r$ converges to a unit vector if and only if one of its components converges to 1. By (12) and (11) this is the case if and only if the following condition (PI') is fulfilled:*

(PI'): *For all $i = 1, \dots, n$ there exists $k \in \{1, \dots, K\}$ such that for all $l \in \{1, \dots, K\} \setminus \{k\}$ there exists $\tilde{s} = \tilde{s}(l) \in \{1, \dots, \hat{s}\}$ such that*

$$c_{l,k}(\tilde{s}) < 0 \quad \text{and} \quad c_{l,k}(s) = 0 \quad \text{for all} \quad s = 1, \dots, \tilde{s} - 1.$$

Checking condition (PI') would give us the limit vector of $\{W_i^{(r)}\}_r$ since it provides the component k , where the vector is 1.

The following corollary shows that in many practical relevant image labeling tasks the iterates in (7) converge the *same* vertex of the simplex.

Corollary 2.6. *Let P be a symmetric stochastic matrix which has eigenvalue 1 of multiplicity 1 and all other eigenvalues have absolute values smaller than 1. Assume that there exists $k \in \{1, \dots, K\}$ such that*

$$\prod_{j=1}^n A_{j,k} > \prod_{j=1}^n A_{j,l} \quad \text{for all} \quad l \in \{1, \dots, K\} \setminus \{k\}. \quad (13)$$

Then, for all $i = 1, \dots, n$, the vectors $W_i^{(r)}$ generated by (7) converge for $r \rightarrow \infty$ to the k -th unit vector.

ε	Labels									
10^{-10}	1	1	3	3	3	3	3	2	2	2
10^{-11}	1	1	1	3	3	3	3	2	2	2
10^{-81}	1	1	1	3	3	3	3	3	2	2
0	3	3	3	3	3	3	3	3	3	3

Table 1. Labeling of a signal with $A = W^{(0)}$ from (14) using different values of ε .

Proof. By assumption on P we have in (PI), resp. (PI'), for $s = 1$ and all $k, l \in \{1, \dots, K\}$ that

$$c_{l,k}(1) = q_{i,1}(\hat{a}_{1,l} - \hat{a}_{1,k})$$

where $\hat{a} = Q^T \log A$. By Lemma 2.1 i) we know that the first column of Q is given by $(q_{i,1})_{i=1}^n = \mathbf{1}_n / \sqrt{n}$. Thus,

$$\hat{a}_{1,m} = \frac{1}{\sqrt{n}} \sum_{j=1}^n \log A_{j,m}.$$

Choosing k according to assumption (13) we conclude $\hat{a}_{1,l} - \hat{a}_{1,k} < 0$ for all $l \in \{1, \dots, K\} \setminus \{k\}$. Thus,

$$c_{l,k}(1) = q_{i,1}(\hat{a}_{1,l} - \hat{a}_{1,k}) = (\hat{a}_{1,l} - \hat{a}_{1,k}) / \sqrt{n} < 0 \quad \text{for all } l \in \{1, \dots, K\} \setminus \{k\}$$

and we can take $\tilde{s} = 1$ in (PI'). Then we see by the proof of Theorem 2.4, resp. Remark 2.5, that $W_i^{(r)}$ converges for $r \rightarrow \infty$ to the k -th unit vector independent of i . \square

By Lemma 2.1 i)-iii), the assumptions on P are fulfilled if P is irreducible and has positive diagonal entries. In our task this is the case if we use local pixel neighborhoods, where the central pixels have positive weights $\rho_{i,i} > 0$, $i = 1, \dots, n$. Assumption (13) is fulfilled in many practical relevant cases, in particular in all our numerical examples. This underlines the relevance of Step 3 in Algorithm 1. We finish this subsection by an intuitive example to illustrate the importance of ε .

Example 2.7. We initialize Algorithm 1 with the matrix

$$A := \begin{pmatrix} \frac{4}{5} & \frac{4}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{10} & \frac{1}{10} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{10} & \frac{1}{10} & \frac{3}{5} & \frac{3}{5} & \frac{3}{5} & \frac{3}{5} & \frac{3}{5} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}^T, \quad (14)$$

which corresponds to a labeling of a signal of length 10 with 3 features and three labels. We choose neighborhoods of size 3 with uniform weights and mirror boundary conditions. Table 1 depicts the results of the labeling for the different ε , where the label assignment changes. In all cases the corresponding label assignment vectors are the vertices of the ε probability simplex (up to machine precision) and the labels 1,2,3 were assigned via (4). Note that we get the same labeling for $\varepsilon = 10^{-1}$ as for $\varepsilon = 10^{-10}$, but the label assignment matrix converges in the first case to vertices as $(0.8, 0.1, 0.1)^T$. Up to $\varepsilon = 10^{-10}$ the pixels are labeled as expected. If ε decreases further, we still get reasonable results originating from the large starting values of the first segment. Until double precision is reached, i.e., $\varepsilon = 10^{-323}$, we

see no more changes in the labeling result. To be able to handle $\varepsilon = 0$ we look at $w = \log(W)$ without normalization, i.e., we iterate (8). Then we observe after 95 iterations that the maximal value of $w^{(95)}$ is attained in the last row. This means that after normalization in exact arithmetic all pixels will get the label 3. Since (13) reads for our example as

$$\left(\prod_{j=1}^n A_{j,k}\right)_{k=1}^3 = \frac{1}{2 \times 10^7} (64, 4, 243),$$

and the third value is the largest one this agrees with Corollary 2.6.

2.3 Gradient Ascent Reprojection Algorithm

In this subsection we show that Algorithm 1 can be seen as gradient ascent reprojection algorithm of a certain function defined on the product manifold of positive numbers. The positive numbers $\mathbb{R}^* := \mathbb{R}_{>0}$ form together with the metric on the tangent space $T_x \mathbb{R}^* = \mathbb{R}$ given by $\langle \xi_1, \xi_2 \rangle_x = \xi_1 \xi_2 / x^2$ a Riemannian manifold. The distance between $x_1, x_2 \in \mathbb{R}^*$ is defined by

$$\text{dist}_{\mathbb{R}^*}(x_1, x_2) := |\log x_1 - \log x_2|.$$

Indeed, computation in \mathbb{R}^* can often be reduced to the Euclidean setting after taking the logarithm. The exponential map $\text{Exp}_x: T_x \mathbb{R}^* \rightarrow \mathbb{R}^*$ and its inverse $\text{Log}_x: \mathbb{R}^* \rightarrow T_x \mathbb{R}^*$ read

$$\text{Exp}_x(\xi) = x \exp \frac{\xi}{x}, \quad \text{Log}_x y = x \log \frac{y}{x}. \quad (15)$$

For $(\gamma_j)_{j=1}^n \in \Delta_n$ and $x_j \in \mathbb{R}^*$, $j = 1, \dots, n$ the *weighted Riemannian center of mass* on \mathbb{R}^* is given by

$$\arg \min_{x \in \mathbb{R}^*} \sum_{j=1}^n \gamma_j \text{dist}_{\mathbb{R}^*}(x, x_j)^2. \quad (16)$$

The minimizer can be computed as follows. The Riemannian gradient of the squared distance is

$$\nabla_{\mathbb{R}^*} \text{dist}_{\mathbb{R}^*}^2(\cdot, y)(x) = -2 \text{Log}_x y = -2x \log \frac{y}{x} = 2 \text{KL}(x, y), \quad (17)$$

where KL denotes the *Kullback-Leibler distance*. It is defined by

$$\text{KL}: \mathbb{R}_{\geq 0}^N \times \mathbb{R}_{>0}^N \rightarrow \mathbb{R}, \quad \text{KL}(x, y) := \langle x, \log \frac{x}{y} \rangle,$$

with the convention $0 \log 0 = 1$ and using the component-wise division in $\frac{x}{y}$. The function is jointly convex in both arguments. It is the Bregman distance of the Shannon entropy and is related to entropy regularized Wasserstein distances, which have recently found applications in image processing, see [44] and the references therein. Setting the Riemannian gradient of the objective in (16) to zero using (17), and dividing by 2 we get

$$0 = \sum_{j=1}^n \gamma_j x \log \frac{x}{x_j} = \sum_{j=1}^n \gamma_j \text{KL}(x, x_j) = \text{KL}\left(x, \prod_{j=1}^n x_j^{\gamma_j}\right). \quad (18)$$

Consequently, the weighted Riemannian center of mass is just the weighted geometric mean $\text{GM} := \prod_{j=1}^n x_j^{\gamma_j}$.

By $\mathcal{M} := (\mathbb{R}^*)^{nK}$ we denote the usual product manifold. We are interested in the function $F: \mathcal{M} \rightarrow \mathbb{R}$ given by

$$\begin{aligned} F(W) &:= \langle \log W, P \log W \rangle \\ &= \sum_{k=1}^K \langle \log W_k, P \log W_k \rangle = \sum_{k=1}^K (\log W_k)^\top P \log W_k, \end{aligned} \quad (19)$$

where W_k , $k = 1, \dots, K$, are the columns of W and P is the weight matrix in (5).

In Step 1 of our algorithm we compute $w = \log W$, perform a gradient ascent step of $f(w) = \langle w, Pw \rangle$ in the Euclidean arithmetic and take the exponential of the result. The direct execution on \mathcal{M} is given by the following lemma.

Lemma 2.8. *Step 1 in Algorithm 1 is a gradient ascent step of F given by (19) at $W^{(r)}$ with step length 1. In other words, we compute*

$$\text{Exp}_W(\nabla_{\mathcal{M}} F(W)) = W \circ \exp(P \log W)$$

at $W = W^{(r)}$, where $\nabla_{\mathcal{M}} F$ denotes the Riemannian gradient of F .

Proof. By separability of F we can restrict our attention to single columns $W = W_k \in (\mathbb{R}^*)^n$ of W . A tangent vector $\xi \in T_W(\mathbb{R}^*)^n$ acts on F as

$$\xi F(W) = \langle \nabla F(W), \xi \rangle,$$

where $\nabla F(W) = \frac{1}{W} \circ P \log W$ is just the Euclidean gradient of F at W . The Riemannian gradient fulfills

$$\langle \nabla_{\mathcal{M}} F(W), \xi \rangle_W = \langle \nabla F(W), \xi \rangle = \left\langle \frac{1}{W} \circ P \log W, \xi \right\rangle$$

for all $\xi \in \mathbb{R}^n$. By the above definition of the Riemannian metric on \mathbb{R}^* we get

$$\langle \nabla_{\mathcal{M}} F(W), \xi \rangle_W = \left\langle \frac{1}{W^2} \circ \nabla_{\mathcal{M}} F(W), \xi \right\rangle = \left\langle \frac{1}{W} \circ P \log W, \xi \right\rangle$$

so that $\nabla_{\mathcal{M}} F(W) = W \circ P \log W$. Using finally (15) gives $\text{Exp}_W(\nabla_{\mathcal{M}} F(W)) = W \circ \exp(P \log W)$. \square

Next we will see that Step 2 and 3 of our algorithm are just projections of the columns of $U^{(r)}$ onto the ε probability simplex defined by

$$\Delta_{K,\varepsilon} := \left\{ x \in \mathbb{R}^K : \sum_{k=1}^K x_k = 1, x \geq \varepsilon \right\} \subset (\mathbb{R}^*)^K, \quad \varepsilon > 0.$$

Here we do not use orthogonal projections, but projections with respect to Kullback-Leibler distance. For the interpretation of the Steps 2-3 in our algorithm as KL projections we need the following lemma.

Lemma 2.9. i) *Assume that $0 < y_1 \leq \dots \leq y_K$. Then, for $0 < \varepsilon < \frac{1}{K}$, the KL projection of y onto $\Delta_{K,\varepsilon}$ is given by*

$$\arg \min_{x \in \Delta_{K,\varepsilon}} \text{KL}(x, y) = (\underbrace{\varepsilon, \dots, \varepsilon}_m, \tau_m y_{m+1}, \dots, \tau_m y_K)^\top,$$

where $\tau_m := \frac{1-m\varepsilon}{\|y\|_1 - \sum_{k=1}^m y_k}$ and m is an index such that

$$y_m \tau_m \leq \varepsilon \quad \text{and} \quad y_{m+1} \tau_m > \varepsilon. \quad (20)$$

ii) For $\varepsilon = 0$ we have

$$\arg \min_{x \in \Delta_K} \text{KL}(x, y) = y / \|y\|_1.$$

Proof. i) We reformulate the convex problem in ii) as

$$\arg \min_{x \in \mathbb{R}^K} \text{KL}(x, y) \quad \text{subject to} \quad \langle x, \mathbf{1}_K \rangle = 1, \quad x - \varepsilon \mathbf{1}_K \geq 0.$$

Using the Lagrangian

$$L(x, p, \tilde{p}) := \text{KL}(x, y) + \tilde{p}(\langle \mathbf{1}_K, x \rangle - 1) + \langle p, x - \varepsilon \mathbf{1}_K \rangle$$

the (necessary and sufficient) optimality conditions for a minimizer read

$$\begin{aligned} (1) \quad & \nabla_x L(x, p, \tilde{p}) = \log \frac{x}{y} + \mathbf{1}_K + \tilde{p} \mathbf{1}_K + p = 0, \\ (2) \quad & \langle x, \mathbf{1}_K \rangle = 1, \quad (3) \quad x - \varepsilon \mathbf{1}_K \geq 0, \\ (4) \quad & p \leq 0, \quad (5) \quad \langle x - \varepsilon \mathbf{1}_K, p \rangle = 0. \end{aligned}$$

The first condition can be rewritten as

$$x = \exp(-\tilde{p} - 1) y \circ \exp(-p)$$

and together with the second condition and $q := \exp(-p)$ we obtain

$$x = \frac{1}{\langle y \circ q, \mathbf{1}_K \rangle} y \circ q = \frac{1}{\langle \frac{y}{\|y\|_1} \circ q, \mathbf{1}_K \rangle} \frac{y}{\|y\|_1} \circ q.$$

By the last equality we can continue our considerations with a vector y fulfilling $\|y\|_1 = 1$.

Assume that $x_1 = \dots = x_m = \varepsilon$ and $x_k > \varepsilon$ for $k = m+1, \dots, K$. Then we have by the fifth condition that $q_k = 1$ for $k = m+1, \dots, K$ so that

$$x = \frac{1}{\sum_{k=1}^m q_k y_k + s_m} (q_1 y_1, \dots, q_m y_m, y_{m+1}, \dots, y_K)^T,$$

where $s_m = \sum_{k=m+1}^K y_k = 1 - \sum_{k=1}^m y_k$. This implies $q_k y_k = c$ for all $k = 1, \dots, m$. Now $x_k = \frac{c}{mc + s_m} = \varepsilon$ gives $c = \frac{s_m \varepsilon}{1 - m\varepsilon} = \frac{\varepsilon}{\tau_m}$ so that

$$q_k = \frac{\varepsilon}{\tau_m y_k}, \quad k = 1, \dots, m.$$

Since the y_k are non-descending, we see that the fourth optimality condition $q_1 \geq \dots \geq q_m = \frac{\varepsilon}{\tau_m y_m} \geq 1$ is fulfilled if the first inequality in (20) holds true. Finally, we have

$$\varepsilon < x_{m+1} = y_{m+1} \tau_m \leq \dots \leq y_K \tau_m = x_K$$

if the second inequality in (20) is satisfied. Thus the vector determined by ii) fulfills all optimality conditions and we are done.

ii) This assertion follows directly from i). \square

As an immediate consequence of the previous lemma and (18) we obtain the following corollary.

Corollary 2.10. *The Steps 2 and 3 in Algorithm 1 are the KL projection of $U_i^{(r)}$, $i = 1, \dots, n$ onto $\Delta_{K,\varepsilon}$. The whole iteration in Step 1-3 can be rewritten as*

$$W_i^{(r+1)} = \arg \min_{W \in \Delta_{K,\varepsilon}} \sum_{j \in \mathcal{N}_\rho(i)} \rho_{i,j} \text{KL}(W, W_i^{(r)} \circ W_j^{(r)}).$$

3 Relation to the Numerical Scheme in [4]

In this section we want to demonstrate how Algorithm 1 is related to an iterative scheme proposed by Åström, Petra, Schmitzer and Schnörr [3, 4, 5] without stretching details on the geometry of the probability simplex which can be found in [4]. The authors consider the interior $\mathring{\Delta}_K$ of the probability simplex (1), which together with the Fisher-Rao metric forms a non complete Riemannian manifold. The tangent space $T_x \mathring{\Delta}_K$ of $\mathring{\Delta}_K$ at $x \in \mathring{\Delta}_K$ is given by the vectors in \mathbb{R}^K whose components sum up to zero. Let $\mathcal{S} := (\mathring{\Delta}_K)^n$ be the corresponding product manifold. By $\text{Exp}_x: T_x \mathcal{S} \rightarrow \mathcal{S}$ we denote the manifold exponential map which, by the incompleteness of the manifold, is only defined on a subset of $T_x \mathcal{S}$. Then they consider the function $E: \mathcal{S} \rightarrow \mathbb{R}$ given by

$$E(W) := \langle S(W), W \rangle,$$

where $S_i(W)$ is the Riemannian center of mass (Karcher mean) of $\{L_j \in \mathring{\Delta}_K : j \in \mathcal{N}(i)\}$ and

$$L(W) := \text{Exp}_W(-U), \quad U_i := \mathcal{D}_i^T - \frac{1}{K} \langle \mathbf{1}_K, \mathcal{D}_i \rangle \mathbf{1}_K$$

with the distance matrix \mathcal{D} given by (2). Since the problem does not admit a closed form solution, the authors propose to apply a Riemannian gradient ascent flow on \mathcal{W} . In their final numerical scheme they utilize certain simplifications, in particular a lifting map for the manifold exponential mapping to cope with the incompleteness of \mathcal{S} . Finally they arrive at the following iterative scheme:

$$W_i^{(0)} := \frac{1}{K} \mathbf{1}_K, \tag{21}$$

$$\begin{aligned} L(W_j^{(r)}) &:= \frac{\exp(-\mathcal{D}_j) \circ W_j^{(r)}}{\langle \exp(-\mathcal{D}_j), W_j^{(r)} \rangle}, \\ S_i(W^{(r)}) &:= \frac{\prod_{j \in \mathcal{N}_\rho(i)} L(W_j^{(r)})^{\rho_{i,j}}}{\langle \mathbf{1}_K, \prod_{j \in \mathcal{N}_\rho(i)} L(W_j^{(r)})^{\rho_{i,j}} \rangle}, \\ W_i^{(r+1)} &:= \frac{W_i^{(r)} \circ S_i(W^{(r)})}{\langle W_i^{(r)}, S_i(W^{(r)}) \rangle}. \end{aligned} \tag{22}$$

where they use the weights $\rho_{i,j} := 1/N_i$, $N_i := |\mathcal{N}_\rho(i)|$.

How does this relate to our setting? Plugging $L(W_j^{(r)})$ into the expression for $S_i(W^{(r)})$ we obtain

$$\begin{aligned} S_i(W^{(r)}) &= \frac{\prod_{j \in \mathcal{N}_\rho(i)} \left(\frac{\exp(-\mathcal{D}_j) \circ W_j^{(r)}}{\langle \exp(-\mathcal{D}_j), W_j^{(r)} \rangle} \right)^{\rho_{i,j}}}{\left\langle \mathbf{1}_K, \prod_{j \in \mathcal{N}_\rho(i)} \left(\frac{\exp(-\mathcal{D}_j) \circ W_j^{(r)}}{\langle \exp(-\mathcal{D}_j), W_j^{(r)} \rangle} \right)^{\rho_{i,j}} \right\rangle} \\ &= \frac{\prod_{j \in \mathcal{N}_\rho(i)} \left(\exp(-\mathcal{D}_j) \circ W_j^{(r)} \right)^{\rho_{i,j}}}{\left\| \prod_{j \in \mathcal{N}_\rho(i)} \left(\exp(-\mathcal{D}_j) \circ W_j^{(r)} \right)^{\rho_{i,j}} \right\|_1} \\ &= \frac{\prod_{j \in \mathcal{N}_\rho(i)} \exp(-\mathcal{D}_j)^{\rho_{i,j}} \circ \prod_{j \in \mathcal{N}_\rho(i)} \left(W_j^{(r)} \right)^{\rho_{i,j}}}{\left\| \prod_{j \in \mathcal{N}_\rho(i)} \exp(-\mathcal{D}_j)^{\rho_{i,j}} \circ \prod_{j \in \mathcal{N}_\rho(i)} \left(W_j^{(r)} \right)^{\rho_{i,j}} \right\|_1}. \end{aligned}$$

Using $A_i := \prod_{j \in \mathcal{N}_\rho(i)} \exp(-\mathcal{D}_j)^{\rho_{i,j}} / \left\| \prod_{j \in \mathcal{N}_\rho(i)} \exp(-\mathcal{D}_j)^{\rho_{i,j}} \right\|_1$, i.e., $\rho_{i,j} = \alpha_{i,j}$ in (3), we get

$$S_i(W^{(r)}) = c A_i \circ \prod_{j \in \mathcal{N}_\rho(i)} \left(W_j^{(r)} \right)^{\rho_{i,j}}, \quad c := \frac{\left\| \prod_{j \in \mathcal{N}_\rho(i)} \exp(-\mathcal{D}_j)^{\rho_{i,j}} \right\|_1}{\left\| \prod_{j \in \mathcal{N}_\rho(i)} \exp(-\mathcal{D}_j)^{\rho_{i,j}} \circ \prod_{j \in \mathcal{N}_\rho(i)} \left(W_j^{(r)} \right)^{\rho_{i,j}} \right\|_1}$$

Substituting this into (22) and dividing the denominator and numerator by the constant c we arrive finally at

$$V_i^{(r+1)} = \frac{A_i \circ W_i^{(r)} \circ \prod_{j \in \mathcal{N}_\rho(i)} (W_j^{(r)})^{\rho_{i,j}}}{\left\| A_i \circ W_i^{(r)} \circ \prod_{j \in \mathcal{N}_\rho(i)} (W_j^{(r)})^{\rho_{i,j}} \right\|_1}. \quad (23)$$

Up to the additional factor A_i formula (23) coincides with the iterations in Algorithm 1.

What about the starting point? Starting with (21) we obtain after the first step in (22) the matrix

$$W_i^{(1)} = \frac{A_i \circ (\mathbf{1}_K / N_i) \circ \prod_{j \in \mathcal{N}_\rho(i)} (\mathbf{1}_K / N_i)^{\rho_{i,j}}}{\left\| A_i \circ (\mathbf{1}_K / N_i) \circ \prod_{j \in \mathcal{N}_\rho(i)} (\mathbf{1}_K / N_i)^{\rho_{i,j}} \right\|_1} = \frac{A_i}{\|A_i\|_1} = A_i.$$

Therefore we can also start with A_i as in Algorithm 1. Clearly, we can use again more flexible weights. In summary the numerical scheme in [4] has the first step

$$U_i^{(r+1)} = A_i \circ W_i^{(r)} \circ \prod_{j \in \mathcal{N}(i)} (W_j^{(r)})^{\rho_{i,j}}. \quad (24)$$

The convergence analysis for $\varepsilon = 0$ is basically the same as in Subsection 2.2. For sake of completeness we add it below. Again we emphasize that that bounding the values away from zero as, e.g., in the third step of our algorithm is not just numerical cosmetics. Without such a step the iterates converge similar as in Corollary 2.6 to the same vertex of the probability simplex which results in a trivial labeling.

Corollary 3.1. *The sequence of iterates $\{W_i^{(r)}\}_r$, $i = 1, \dots, n$ generated by (24) converges as $r \rightarrow \infty$. It converges to a unit vector if and only if A fulfills (PI) but with*

$$c_{l,k}(s) := \sum_{j=n_s}^{n_s+\kappa_s-1} q_{i,j} \lambda_j^{-1} (\hat{a}_{j,l} - \hat{a}_{j,k}).$$

Proof. The proof follows the same lines as those of Theorem 2.4 so that we highlight only the differences. Instead of (10) we obtain

$$\begin{aligned} w^{(r)} &= (I + (I + P) + \dots + (I + P)^r) w^{(0)} = P^{-1} ((I + P)^{r+1} - I) w^{(0)} \\ &= Q \Lambda^{-1} (D^{r+1} - I) \hat{a}. \end{aligned}$$

Taking the componentwise exponential and normalizing we get

$$W_i^{(r)} = \left((1 + \sum_{\substack{l=1 \\ l \neq k}}^K G_{l,k}(r))^{-1} \right)_{k=1}^K,$$

where

$$G_{l,k}(r) := \frac{\exp(q_i^T \Lambda^{-1} (D^{r+1} - I) \hat{a}_l)}{\exp(q_i^T \Lambda^{-1} (D^{r+1} - I) \hat{a}_k)}.$$

In particular we have

$$\begin{aligned} G_{l,k}(r) &= \prod_{j=1}^n \exp \left(q_{i,j} (\mu_j^{r+1} - 1) \lambda_j^{-1} (\hat{a}_{j,l} - \hat{a}_{j,k}) \right), \\ g_{l,k}(r) &= \log(G_{l,k}(r)) = \sum_{j=1}^n q_{i,j} (\mu_j^{r+1} - 1) \lambda_j^{-1} (\hat{a}_{j,l} - \hat{a}_{j,k}) \\ &= \sum_{s=1}^M c_{l,k}(s) (\mu_{n_s}^{r+1} - 1). \end{aligned}$$

The final conclusions are analogous as in the proof of Theorem 2.4. \square

4 Numerical Examples

In this section we demonstrate the performance of Algorithm 1 for the partitioning of images. We start with color images having values in the Euclidean space and turn to manifold-valued images afterwards. In particular we consider diffusion tensor magnetic resonance images (DT-MRI) with values in the Hadamard manifold of symmetric positive definite matrices and electron backscatter diffraction (EBSD) images, whose values are taken from quotient spaces of the manifold of rotations $\text{SO}(3)$.

Let $\mathcal{G} := \{1, \dots, N\} \times \{1, \dots, M\}$ be the image grid and $\emptyset \neq \mathcal{V} \subseteq \mathcal{G}$ the set of available, in general noisy pixels. In particular, we have $\mathcal{V} = \mathcal{G}$ in the case of no missing pixels. We are interested in labeling an image $f: \mathcal{V} \rightarrow \mathcal{M}$.

In our numerical examples we stop if the average entropy of the labeling matrix given by

$$-\frac{1}{|\mathcal{G}|} \sum_{i \in \mathcal{G}} \sum_{k=1}^K W_{i,k}^{(r)} \log(W_{i,k}^{(r)}),$$

drops below a certain threshold. If not stated otherwise, we take 10^{-3} as threshold as in [4], where the same stopping criterion was suggested. If not stated otherwise the same weights were used for α and ρ .

The algorithm was implemented in MATLAB 2016a.

Color Images. In our *first example* we label the `mandrill` image $f: \mathcal{G} \rightarrow [0, 1]^3 \subset \mathbb{R}^3$ in Fig. 2 (a), which was also handled in [4]. We use the 20 color labels shown in Fig. 2 (b) as prior features which were picked manually from the original image. We mention that the numerical results with Step 1 as in (24), i.e., with the additional factor A_i proposed in [4] differ only in 4 percent of the image pixels from those presented below so that there is no visual difference.

First, we are interested in the influence of different weights α, ρ , namely

- i) *uniform local* weights $\rho_{i,k} := 1/s^2$, where $s^2 := |\mathcal{N}(i)|$ is the size of the local $s \times s$ pixel neighborhood.
- ii) *nonlocal* weights which are constructed as follows: a similarity measure for two pixels $i, j \in \mathcal{G}$ is given by the weighted difference of the surrounding $s_p \times s_p$ patches

$$d_p(i, j) := \sum_{l_1, l_2 = -s_p}^{s_p} g_{\sigma_p}(l_1, l_2) \text{dist}(f_{i+(l_1, l_2)^T}, f_{j+(l_1, l_2)^T}),$$

where $g_{\sigma_p}: \mathbb{R}^2 \rightarrow \mathbb{R}$ is a Gaussian with standard deviation σ_p . The neighborhood $\tilde{\mathcal{N}}(i)$ is then given by the s_{nl} pixels with smallest distances d_p to i . Then we make the neighborhood symmetric by setting

$$\mathcal{N}_{nl}(i) := \tilde{\mathcal{N}}(i) \cup \{j \in \mathcal{G} | i \in \tilde{\mathcal{N}}(j)\}.$$


To reduce the computational effort for computing the weights we just compare each pixel with the pixels inside a $t \times t$ neighborhood. Finally the weights $\rho_{i,j}$ are defined as

$$\rho_{i,j} := \frac{\tilde{\rho}_{i,j}}{\sum_{j \in \mathcal{N}_{nl}(i)} \tilde{\rho}_{i,j}}, \quad \tilde{\rho}_{i,j} := \exp\left(-\frac{d_p(i, j)}{2\sigma_w^2}\right).$$


We mirror the image at the boundary.

Fig. 2 shows the labeling with Algorithm 1, more precisely:


- α *local*, ρ *local*: In Figs. 2 (c)–(e) we used uniform local weights both for α and ρ of different sizes. We see that large neighborhoods lead to coarse label assignments. The number of iterations needed to reach the stopping criterion also increases with the size of the neighborhood from 27 to 39.




(a) Original image mandrill



(b) Prior features



(c) Local weights, $s = 3$



(d) Local weights, $s = 5$



(e) Local weights, $s = 7$



(f) Nonlocal weights,
 $s_{nl} = 7$




(g) Nonlocal weights,
 $s_{nl} = 19$




(h) Nonlocal weights,
 $s_{nl} = 37$



(i) α from (c), ρ from (f)



(j) α from (d), ρ from (g)



(k) α from (e), ρ from (h)

Figure 2. Illustration of the influence of the chosen neighborhood in using Algorithm 1 for labeling the mandrill with 20 different labels and .

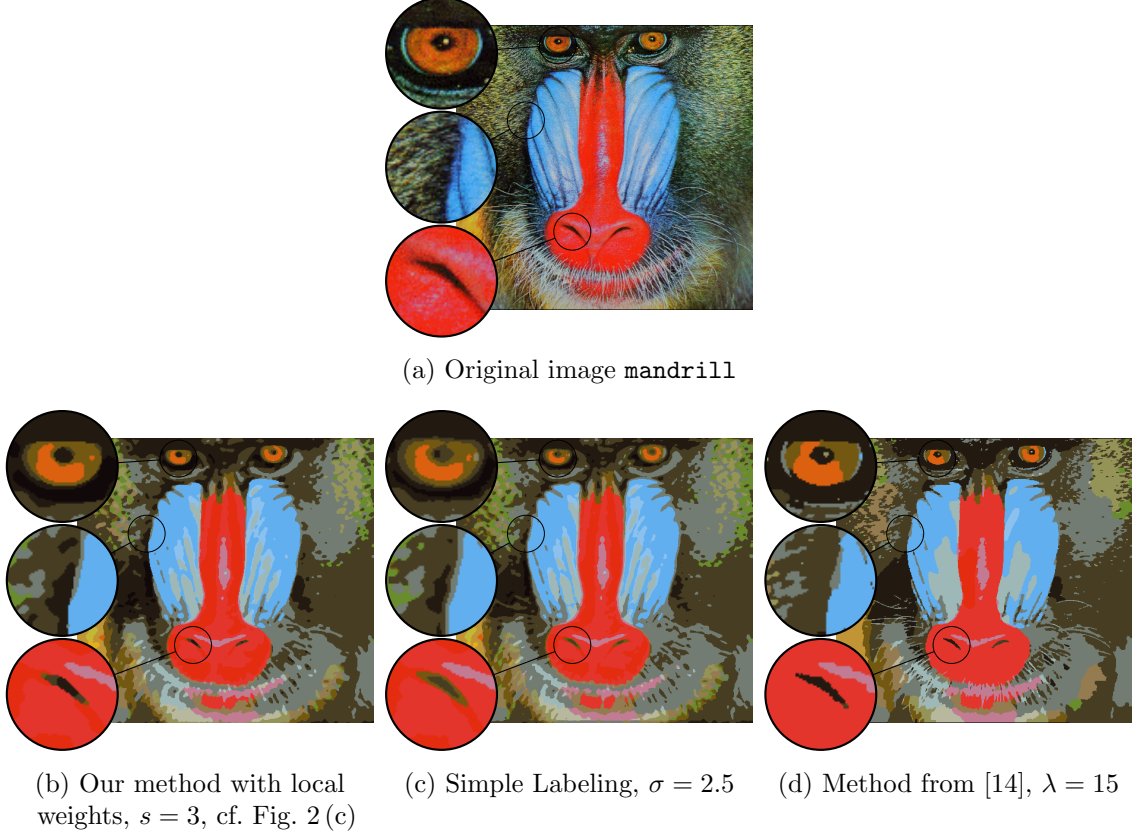


Figure 3. Comparison of (b) our Algorithm 1 with two other labeling methods: (c) a Simple Labeling and (d) a variational model from [14] using the labels from Fig. 2 (b).

- α *nonlocal*, ρ *nonlocal*: Using nonlocal neighborhoods with parameters $s_p = 3$, $s_{nl} \in \{7, 19, 37\}$, $t = 11$, $\sigma_p = 1$, $\sigma_w = 0.2$ both for α and ρ results in Figs. 2 (f)–(h). Even though we take a similar number of neighbors in each column of Fig. 2, we see that the nonlocal approach preserves the most details of the original image, such as the pupils or the whiskers.
- α *local*, ρ *nonlocal*: In Figs. 2 (i)–(k), we look at the influence of choosing different α and ρ . Here we choose α as uniform local weights and ρ as nonlocal weights from the respective column of Fig. 2. The results are a mixture of the other two label assignments. On the one hand we obtain larger connected regions as the purely nonlocal approach on the other hand we keep finer details such the pupils and nostrils in comparison to the local approach.
- α *nonlocal*, ρ *local*: Switching α and ρ leads to results very similar to those shown in Figs. 2 (c)–(e).

Next we want to compare the proposed algorithm with two other labeling methods. The first method, called *Simple Labeling* in the following, starts by computing the convolution of each of the RGB channels with a Gaussian of standard deviation σ . The label assignment is then given by choosing the closest prior feature at each pixel, i.e., the label of pixel i is given

by $\arg \min_k \text{dist}(f_{\sigma,i}, f_k^*)$, where f_σ is the image f convolved with the Gaussian. Note that for more general manifold-valued images as considered in the next examples, the Riemannian center of mass would be the natural choice for averaging, but its computation involves a minimization process and can not be stated explicitly. The second method we compare with is a variational method given in [14]. Of course there exists many more methods to compare to; we chose these two, because the first one is even simpler, and the second one is a state-of-the-art variational model for image partitioning.

In Fig. 3 we compare these two methods with the multiplicative labeling derived in this paper. In order to compare the methods, we choose the parameters s, σ, λ from our approach, the Simple Labeling and the method from [14], respectively, such that they result in similar labeling in the background, e.g., the fur in the middle right part of the image of the **mandrill**. We further focus on three regions: the left eye, the left side of the nose and the left nostril. The original image, shown in Fig. 3(a), is labeled using the prior features from the last example, cf. Fig. 2(b).

When we compare our method with local weights and $s = 3$, cf. Fig. 2(c), whose focus regions are shown in Fig. 2(c) to the result of the Simple Labeling, cf. Fig. 3(c), $\sigma = 2.5$, the details are kept sharper in our method, while they get blurred in the latter approach. This can clearly be seen at the nostril or the nose detail, where at least one additional label is introduced at the boundary, and the pupil is kept in its original size in our approach. However, in e.g. more homogeneous regions like the red upper part of the nose or the fur in the top right, both methods result in similar label assignments.

Compared to the variational method from [14] our approach seems to keep a little more detail e.g. at the boundary of the nose or the fur in the bottom right, while just losing a little bit of detail, like the white spot in the eye. Despite of these two minor differences they perform similar.

In our *second example* we apply the proposed Algorithm 1 using local weights as in i) in order to label the **pepper** image, see Fig. 4(a). At this point we assume that we are only given data at grid points $\mathcal{V} \subset \mathcal{G}$, $|\mathcal{V}| \ll |\mathcal{G}|$, where 95% of the pixel values are unknown, cf. Fig. 4(b). The prior features are chosen by a discretization of the color cube, more precisely we consider, $\{(f_1^*, f_2^*, f_3^*) : f_i^* \in \{0, \frac{1}{7}, \dots, 1\}, i \in \{1, 2, 3\}\}$. The distance matrix $\mathcal{D} := (d_{i,k})_{i=1,k=1}^{n,K}$ is given by

$$d_{i,k} := \begin{cases} \text{dist}(f_i, f_k^*) & i \in \mathcal{V}, \\ 0 & i \notin \mathcal{V}. \end{cases}$$

Fig. 4(d) shows the resulting labeling. Comparing this with the labeling of the original image shown in Fig. 4(c) we see that our approach can cope with missing data very well.


Symmetric Positive Definite Matrices. We want to label images having values on the manifold of symmetric positive definite $r \times r$ matrices $\mathcal{P}(r)$ with the affine invariant distance, see e.g. [2, 42],

$$\text{dist}_{\mathcal{P}(r)}(x_1, x_2) := \|\text{Log}(x_1^{-\frac{1}{2}} x_2 x_1^{-\frac{1}{2}})\|,$$

where $\text{Log } x := -\sum_{k=1}^{\infty} \frac{1}{k} (I - x)^k$, $\rho(I - x) < 1$, denotes the matrix logarithm of $x \in \mathcal{P}(r)$ and the norm is the Frobenius norm.



(a) Original image peppers



(b) Lossy image, 95% of missing pixels



(c) Labeling of the original image



(d) Labeling of the lossy image

Figure 4. Comparison between labeling the original image (a) and the lossy one (b) for uniform local weights with $s = 5$.

As a *first example* we look at diffusion tensors occurring in magnetic resonance imaging (DT-MRI). For DT-MRI a 3D dataset from MRI is used to compute diffusion tensors at each measurement position resulting in a dataset, where each data point is a symmetric positive definite matrix. Hence all data items are given on the manifold $\mathcal{P}(3)$ of symmetric positive definite 3×3 matrices. The Camino project¹ [18] provides a freely available dataset of DT-MRI of the human head. We take a look at the traversal cut at $z = 35$ in the dataset, i.e., our data is a manifold-valued image of size $N \times M$ with $N = M = 112$ where the pixels outside the head are missing, see Fig. 5 (a). The coloring is based on the anisotropy index, cf. [37], employing the Matlab colormap `hsv`.

We ignore the outer pixels in the following by setting their distance (and hence influence) to any other pixel to zero. We define two prior features, which are depicted as pink ellipses, one denoting the inner, less diffusion affected area, one to capture the boundary and main diffusion areas of the brain. We compare three approaches, i) an equally weighted neighborhood of size $s \times s$, ii) a weighted 3×3 using a truncated Gaussian, $\sigma = 0.5$ and iii) the nonlocal weights. We set $s_p = 3$, $s_{nl} = 9$, $t = 15$, and for the two involved Gaussians $\sigma_p = 3$ and $\sigma_w = \frac{1}{5}$. The results are shown in Figs. 5 (b)–(d): while the equally weighted neighborhood captures the boundary of the brain correctly despite a small part at the top right, it is not able to “follow” the diffusivity paths inside the brain. Using the weighted neighborhood captures the main features while loosing a little detail on the top right border. The nonlocal approach captures the complete boundary and even more details of diffusion paths inside the brain.

In our *second example* for symmetric positive definite matrices we have a look at the texture segmentation approach from [49]. There the authors assign to each pixel of the image a covariance matrix, which is symmetric positive semidefinite by definition, and we can assume it is positive definite for a large enough number of samples in natural images. These covariance matrices are based on the following feature vectors

$$F(i) = (I(i), I_x(i), I_y(i), I_{xx}(i), I_{yy}(i))^T,$$

where I is the intensity of the image, I_x, I_{xx} , and I_y, I_{yy} , are the first and second derivatives in horizontal, respectively vertical direction. To avoid the influence of noise we convolve the image with a Gaussian of small standard deviation σ before computing the feature vectors. The derivatives are discretized via central differences. Given a neighborhood $\mathcal{N}_{\text{cov}}(i)$ around each pixel the covariance matrices are defined as


$$C(i) = \frac{1}{|\mathcal{N}_{\text{cov}}(i)| - 1} \sum_{j \in \mathcal{N}_{\text{cov}}(i)} (F(j) - \mu(i))(F(j) - \mu(i))^T,$$

where $\mu(i) \in \mathbb{R}^5$ is the mean of the features F in the neighborhood $\mathcal{N}_{\text{cov}}(i)$. To achieve a good segmentation result we also include the mean μ into our final features. So we work on the product manifold $\mathbb{R}^5 \times \mathcal{P}(5)$ with distance


$$\text{dist}_{\mathbb{R}^5 \times \mathcal{P}(5)}((\mu_1, C_1), (\mu_2, C_2)) = \sqrt{\|\mu_1 - \mu_2\|_2^2 + \text{dist}_{\mathcal{P}(5)}^2(C_1, C_2)}.$$

If we want to segment the texture image in Fig. 1 (a) we need appropriate prior features. They are obtained from six supervising texture images using 100 randomly selected rectangular


¹see <http://camino.cs.ucl.ac.uk>




(a) Original DT-MR image



(b) Equal weights, $s = 3$



(c) Gaussian weights on 3×3 neighborhood, $\sigma = 0.5$



(d) Nonlocal weights

Figure 5. Segmentation of DT-MRI data with two labels and different neighborhoods.

patches and calculate the mean of their means and covariance matrices. The result in Fig. 1 (b) is then obtained by using $\sigma = 0.5$ for the convolution, 7×7 patches as neighborhoods $\mathcal{N}_{\text{cov}}(i)$ and 15×15 patches as neighborhoods \mathcal{N}_ρ with uniform weights in Algorithm 1.

EBSD Data. The analysis of polycrystalline materials by means of electron backscattered diffraction (EBSD) is a fundamental task in material science, see [1, 33]. Since the microscopic grain structure affects macroscopic attributes of materials such as ductility, electrical and lifetime properties, there is a growing interest in the grain structure of crystalline materials such as metals and minerals. EBSD provides for each position on the surface of a specimen a so-called Kikuchi pattern, which allows the identification of

- i) the structure (material index), and
- ii) the orientation of the crystal relative to a fixed coordinate system ($\text{SO}(3)$ value).

Since the atomic structure of a crystal is invariant under its specific finite symmetry group $S \subset \text{SO}(3)$, the orientation at each grid point $i \in \mathcal{G}$ is only given as an equivalence class

$$[f_i] = \{f_i s \mid s \in S\} \in \text{SO}(3)/S, \quad f_i \in \text{SO}(3).$$

The following remark introduces the colorization method which is implemented in the software package MTEX [6, 39] and which we will use for visualization.

Remark 4.1. *The colorization of $\text{SO}(3)/S$ valued rotations $[f]$ is done as follows: for a fixed vector $\vec{r} \in \mathbb{S}^2$ we consider the mapping $\Phi: \text{SO}(3)/S \rightarrow \mathbb{S}^2/S$, $[f] \mapsto [f^{-1}\vec{r}]$ and then we assign a certain color to each point on \mathbb{S}^2/S . Usually \vec{r} is defined as the vector orthogonal to the specimen surface. Note that this reduces the dimension by one since we can no longer distinguish between rotations around \vec{r} . After applying Φ we introduce a color coding of the sphere that takes into account the symmetries.*

To illustrate the color coding, we choose quartz as an example as it has neither trivial nor too complicated symmetries. The colorization of the quotient \mathbb{S}^2/S for this example is depicted in Fig. 6. The color map shown in Fig. 6 c) respects the symmetry group $S \subset \text{SO}(3)$ of quartz. The symmetry group has three rotations with respect to a 3-folded axis ($2k\pi/3$, $k = 1, \dots, 3$ rotations around the main axis c of the crystal) and three rotations with respect to 2-folded axis, e.g., a_1 , a_2 , perpendicular to c . The polyhedron drawn in Fig. 6 a) and b) to visualize the orientation has exactly these symmetries.

The grains in EBSD images are regions with similar orientations which we want to label. As distance between orientations on $\text{SO}(3)$ we use the geodesic distance

$$\text{dist}_{\text{SO}(3)}(f_1, f_2) := \arccos|\langle f_1, f_2 \rangle|,$$

where we suppose that the orientations are given as quaternions, see [24], i.e., $f_i \in \mathbb{S}^3/\{-1, 1\}$. To cope with the symmetry, we apply the usual distance on the quotient space $\text{SO}(3)/S$, cf. [21, p. 153]:

$$\text{dist}_{\text{SO}(3)/S}(f_1, f_2) := \min_{\tilde{f}_1 \in [f_1]_S, \tilde{f}_2 \in [f_2]_S} \text{dist}_{\text{SO}(3)}(\tilde{f}_1, \tilde{f}_2).$$




Figure 6. Visualization of the rotations considering the symmetry group of quartz according to the mapping $\text{SO}(3)/S \ni [f] \mapsto [f^{-1}r] \in \mathbb{S}^2/S$.

In EBSD data usually pixels are missing, e.g., due to corrupted Kikuchi patterns. Therefore we extend the approach to this setting and modify the distance as in the previous paragraph: if a pixel is missing, the distance between this pixel and every cluster center is set to 0. As prior features we randomly choose orientations in the original data in such a way that the distance between all prior features is at least 0.3 and for each pixel in the original data there exists a prior feature with distance less or equal to 0.3.

In our *first example* we label EBSD data obtained from the single phase material Magnesium, where few pixels are missing, see [6]. We focus on Algorithm 1. We fix $\nu = 1$ and choose Gaussian weights $\alpha_{i,k} = \rho_{i,k}$ with standard deviation $\sigma = 0.8$ which are truncated outside $[-3\sigma, 3\sigma]$, i.e., we work in a local 5×5 neighborhood. Again, we mirror the image at the boundary. The result is shown in Fig. 7. In contrast to the method for finding grain boundaries via thresholding described in [7] and implemented in the software package MTEX [6] the proposed approach yields smoother grain boundaries and is able to reduce noise. The grain boundaries are sometimes too fringy in the original data. This happens since the diffraction patterns are inaccurate at grain boundaries due to the change in orientation and therefore also the assigned rotations are inaccurate. Here, the labeling method performs better.

In the *second example* we draw our attention to multiphase EBSD data. We consider a multiphase dataset consisting of the three phases Forsterite, Enstatite and Diopside taken from MTEX [6]. These phases have different crystal symmetries such that we can not directly compute the distance between orientations belonging to different phases. Therefore, the features have to represent not only the orientations, but also the phase they are corresponding to. We consider features (f_i, p_i) where $p_i \in \{1, 2, 3\}$ denotes the phase and $f_i \in \text{SO}(3)/S(p_i)$ is the orientation where $S(p_i)$ denotes the symmetry group of the phase p_i . Then we choose a finite, but constant distance τ between pixels that belong to different phases to allow for




Figure 7. Labeling of single phase EBSD data (Magnesium) with 31 prior features. \mathcal{N} is chosen as the 5×5 neighborhood with truncated Gaussian weights ($\sigma = 0.8$). The proposed model is able to reduce noise and produces smoother boundaries.

phase changes of noisy pixels. Now, the distance matrix $\mathcal{D} := (d_{i,k})_{i=1,k=1}^{n,K}$ has entries

$$d_{i,k} := \begin{cases} 0 & i \notin \mathcal{V}, \\ \text{dist}_{\text{SO}(3)/S(p_i)}(f_i, f_k^*) & i \in \mathcal{V}, p_i = p_k^*, \\ \tau & \text{otherwise.} \end{cases}$$

Fig. 8 shows an example of such a data set. Here we set $\tau := 1$. We see that the partitioning of both, phase and orientation works well even in the presence of missing (white) pixels. In particular, also the phase boundaries get smoothed. Furthermore, noise in the phase data, i.e., pixels that were assigned to the wrong phase in the image acquisition process, is reduced.

5 Conclusions

We have proposed a simple and efficient algorithm for data labeling based on an iterative geometric averaging of a label assignment matrix. The algorithm starts with a matrix of averaged distances between the prior features and those which we want to label. In applications this distance may not be the Euclidean one, but e.g., the distance on an appropriate manifold. Since the weights in the label assignment matrix are updated using only few of their local or nonlocal neighbors, our algorithm has a high potential for parallelization, which is not exploited in the implementation so far. In particular, a large number of labels does not increase the convergence speed drastically, which is the case for various TV-regularized variational methods as [8, 14, 17, 16, 36].

There are still a lot of open questions, which we want to examine in our future work, e.g.:

- i) the convergence analysis of Algorithm 1 with Step 3, in particular the role of ε . Note that we have observed convergence to the vertices of $\Delta_{K,\varepsilon}$ in all our numerical examples.

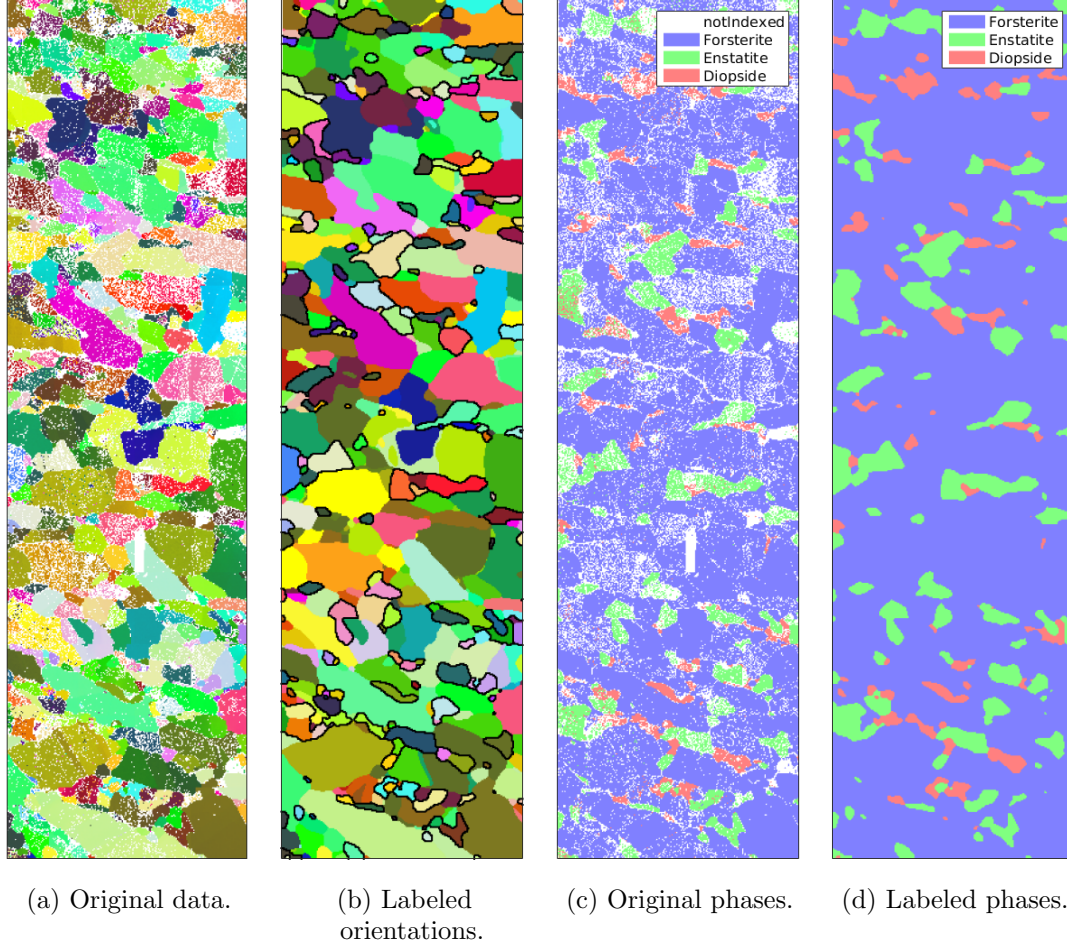


Figure 8. Labeling of corrupted three-phase EBSD data with 447 prior orientations by Algorithm 1. For the original data also the MTEX colorization was applied. Phase transitions are emphasized by black lines in the labeled orientations.

ii) and to the solution of variational problems.

Of course it would be useful to have a similar multiplicative filtering approach for non supervised labeling or to develop updating strategies for the prior features. Further, we intend to use the method for other applications.

Acknowledgement. We are grateful to Ch. Schnörr (University of Heidelberg) for stimulating discussions. Many thanks to R. Hielscher (University of Chemnitz) for supporting the work on EBSD data. We thank the referees for requesting a discussion of condition (PI). Funding by the German Research Foundation (DFG) within the project STE 571/13-1 & BE 5888/2-1 and within the Research Training Group 1932 “Stochastic Models for Innovations in the Engineering Sciences”, project area P3, is gratefully acknowledged. Furthermore, G. Steidl acknowledges the support by the German Federal Ministry of Education and Research (BMBF) through grant 05M13UKA (AniS).

References

- [1] B. L. Adams, S. I. Wright, and K. Kunze. Orientation imaging: The emergence of a new microscopy. *Journal Metallurgical and Materials Transactions A*, 24:819–831, 1993.
- [2] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, 2006.
- [3] F. Åström, S. Petra, B. Schmitzer, and C. Schnörr. A Geometric Approach to Image Labeling. In *Proc. ECCV*, 2016.
- [4] F. Åström, S. Petra, B. Schmitzer, and C. Schnörr. Image labeling by assignment. *ArXiv Preprint 1603.05285*, 2016.
- [5] F. Åström, S. Petra, B. Schmitzer, and C. Schnörr. The Assignment Manifold: A Smooth Model for Image Labeling. In *Proc. 2nd Int. Workshop on Differential Geometry in Computer Vision and Machine Learning*, 2016.
- [6] F. Bachmann and R. Hielscher. MTEX – MATLAB toolbox for quantitative texture analysis. <http://mtex-toolbox.github.io/>, 2005–2016.
- [7] F. Bachmann, R. Hielscher, and H. Schaeben. Grain detection from 2d and 3d EBSD data—specification of the MTEX algorithm. *Ultramicroscopy*, 111(12):1720–1733, 2011.
- [8] E. Bae, J. Yuan, and X.-C. Tai. Global minimization for continuous multiphase partitioning problems using a dual approach. *International Journal of Computer Vision*, 92(1):112–129, 2011.
- [9] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2003.
- [10] A. Buades, B. Coll, and J.-M. Morel. Neighborhood filters and PDEs. *Numerische Mathematik*, 105:1–34, 2006.

- [11] A. Buades, B. Coll, and J.-M. Morel. Image denoising methods. A new nonlocal principle. *SIAM Review*, 52(1):113–147, 2010.
- [12] M. Burger, A. Sawatzky, and G. Steidl. First order algorithms in variational image processing. In R. Glowinski, S. Osher, and W. Yin, editors, *Operator Splittings and Alternating Direction Methods*. Springer, 2016.
- [13] X. Cai, R. Chan, and T. Zeng. A two-stage image segmentation method using a convex variant of the Mumford–Shah model and thresholding. *SIAM Journal on Imaging Sciences*, 6(1):368–390, 2013.
- [14] A. Chambolle, D. Cremers, and T. Pock. A convex approach to minimal partitions. *SIAM Journal on Imaging Sciences*, 5(4):1113–1158, 2012.
- [15] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [16] T. F. Chan, S. Esedoglu, and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal on Applied Mathematics*, 66(5):1632–1648, 2006.
- [17] C. Chaux, A. Jezierska, J.-C. Pesquet, and H. Talbot. A spatial regularization approach for vector quantization. *Journal of Mathematical Imaging and Vision*, 41(1-2):23–38, 2011.
- [18] P. A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, page 2759, Seattle, WA, USA, 2006.
- [19] C. A. Deledalle, L. Denis, and F. Tupin. Iterative weighted maximum likelihood denoising with probabilistic patch-based weights. *IEEE Transactions on Image Processing*, 18(12):2661–2672, 2009.
- [20] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing*, 15(12):3736–3745, 2006.
- [21] G. B. Folland. *Real Analysis*. Wiley & Sons, 1999.
- [22] G. F. Frobenius. *Über Matrizen aus nichtnegativen Elementen*. Königliche Akademie der Wissenschaften, 1912.
- [23] G. Gilboa and S. Osher. Nonlocal linear image regularization and supervised segmentation. *SIAM Journal on Multiscale Modeling and Simulation*, 6(2):595–630, 2007.
- [24] M. Gräf. A unified approach to scattered data approximation on \mathbb{S}^3 and $\text{SO}(3)$. *Advances in Computational Mathematics*, 37:379–392, 2012.
- [25] S. Häuser and G. Steidl. Convex multiclass segmentation with shearlet regularization. *International Journal of Computer Mathematics*, 90(1):62–81, 2013.

- [26] L. Herault and R. Horaud. Figure-ground discrimination: A combinatorial optimization approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):899–914, 1993.
- [27] T. Hofmann and J. M. Buhmann. Pairwise data clustering by deterministic annealing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(1):1–14, 1997.
- [28] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1986.
- [29] R. A. Hummel and S. W. Zucker. On the foundation of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5(3):267–287, 1983.
- [30] J. H. Kappes, B. Andres, F. A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, T. Kröger, J. Lellmann, et al. A comparative study of modern inference techniques for structured discrete energy minimization problems. *International Journal of Computer Vision*, 115(2):155–184, 2015.
- [31] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1568–1583, 2006.
- [32] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, 2004.
- [33] K. Kunze, S. I. Wright, B. L. Adams, and D. J. Dingley. Advances in automatic EBSD single orientation measurements. *Textures and Microstructures*, 20:41–54, 1993.
- [34] F. Laus, J. Persch, and G. Steidl. A nonlocal denoising algorithm for manifold-valued images using second order statistics. *ArXiv Preprint 1607.08481*, 2016.
- [35] J. Lellmann, F. Lenzen, and C. Schnörr. Optimality bounds for a variational relaxation of the image partitioning problem. *Journal of Mathematical Imaging and Vision*, 47(3):239–257, 2013.
- [36] J. Lellmann and C. Schnörr. Continuous multiclass labeling approaches and algorithms. *SIAM Journal on Imaging Sciences*, 4(4):1049–1096, 2011.
- [37] M. Moakher and P. G. Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In *Visualization and Processing of Tensor Fields*, pages 285–298. Springer, Berlin, Heidelberg, 2006.
- [38] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42(5):577–685, 1989.
- [39] G. Nolze and R. Hielscher. IPF coloring of crystal orientation data. *Preprint Technische Universität Chemnitz*, 2016.

- [40] H. Orland. Mean-field theory for optimization problems. *Journal de Physique Lettres*, 46(17):763–770, 1985.
- [41] M. Pelillo. The dynamics of nonlinear relaxation labeling processes. *Journal of Mathematical Imaging and Vision*, 7:309–323, 1997.
- [42] X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66:41–66, 2006.
- [43] O. Perron. Zur Theorie der Matrizen. *Mathematische Annalen*, 64(2):248–263, 1907.
- [44] G. Peyré. Entropic Wasserstein gradient flows. *ArXiv Preprint 1502.06216v3*, 2015.
- [45] A. Rosenfeld, R. A. Hummel, and S. W. Zucker. Scene labeling by relaxation operations. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-6(6):420–433, 1976.
- [46] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1):259–268, 1992.
- [47] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. A comparative study of energy minimizing methods for Markov random fields with smoothness-based priors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(6):1068–1080, 2008.
- [48] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. Sixth International Conference on Computer Vision*, pages 839–846, Bombay, India, Jan. 1998. Narosa Publishing House.
- [49] O. Tuzel, F. Porikli, and P. Meer. Region covariance: A fast descriptor for detection and classification. In *European Conference on Computer Vision*, pages 589–600. Springer, 2006.
- [50] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- [51] J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998.
- [52] H. Wielandt. Unzerlegbare, nicht negative matrizen. *Mathematische Zeitschrift*, 52(1):642–648, 1950.
- [53] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51(7):2282–2312, 2005.